

AD-A234 753

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
<small>Public reporting burden for this report is estimated to be 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, U.S. Government Printing Office, Washington, DC 20540-6001.</small>				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE	3. REPORT TYPE AND DATES COVERED	
			FINAL 01 Jul 90 to 31 Dec 90	
4. TITLE AND SUBTITLE			5. FUNDING NUMBERS	
INTEGRATED DISTRIBUTED LOGISTICS MANAGEMENT			F49620-90-C-0035	
6. AUTHOR(S)			65520F 3005/A1	
Dr. Nutly				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)			8. PERFORMING ORGANIZATION REPORT NUMBER	
Computer Aided Planning & Scheduling Inc. 2900 Paces Ferry Road, Building D Atlanta, GA 303393-3719			AFOSR-TR- 91 0241	
9. SPONSORING MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSORING MONITORING AGENCY REPORT NUMBER	
AFOSR/HM Bldg 410 Bolling AFB DC 20332-6448			F49620-90-C-0035	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT			12b. DISTRIBUTION CODE	
Approved for public release; distribution unlimited.				
13. ABSTRACT (Maximum 200 words)				
Contractor has devised a new method of solution for problems which are difficult due to their size & multi-level nature. Different objectives, constraints and techniques are used at different levels of the hierarchy.				
14. SUBJECT TERMS			15. NUMBER OF PAGES	
91 4 16 037				
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	
UNCLASSIFIED	UNCLASSIFIED	UNCLASSIFIED	UL	

INTEGRATED DISTRIBUTED LOGISTICS MANAGEMENT

SBIR PHASE I (F49620-90-C-0035)

FINAL TECHNICAL REPORT

February 28, 1991

Prepared by

COMPUTER AIDED PLANNING & SCHEDULING, INC.

2900 Paces Ferry Road, Building D

Atlanta, Georgia 30339-3719

11/19/90 11:11:30

SEARCHED	INDEXED
SERIALIZED	FILED
JUN 1991	
FBI - ATLANTA	
A-1	

TABLE OF CONTENTS

1. INTRODUCTION
2. PHASE I RESEARCH OBJECTIVES
3. PHASE I ACCOMPLISHMENTS
 - 3.1 Prototype Problem Scenario
 - 3.2 Distributed Logistics Management Approach
 - 3.3 Model of Distributed Logistics Management
4. PROTOTYPE IMPLEMENTATION AND FUNCTIONAL DESCRIPTION
 - 4.1 Overview
 - 4.2 Low-Level Tactical/Operational Planning Modules
 - 4.3 High-Level Tactical Planning Module
 - 4.4 Information Management System
5. RELATIONSHIP TO FUTURE RESEARCH
6. CONCLUSION

1. INTRODUCTION

Atlanta-based Computer Aided Planning & Scheduling, Inc. (CAPS) is nearing completion of a Phase I Small Business Innovation Research (SBIR) award from the Air Force Office Of Scientific Research entitled "Interactive Distributed Logistics Management." The principal idea is to develop computer-aided methodology to allow different logistics planning levels to simultaneously and cooperatively build and administer a logistics plan. This methodology will allow the military and other organizations to use integrated computer-aided models in a manner closely matching the organizational decision hierarchy, resulting in higher quality plans developed faster. Operation Desert Shield in particular has spotlighted the importance and complexity of successfully coordinating large-scale logistics planning.

Prior to the 1980's, computer aided models for military deployment planning were primarily based on simulation concepts. These models executed in a batch fashion on mainframe-size computers and were basically descriptive, playing out an existing logistics plan in order to examine its behavior and feasibility. In the late 1970's, researchers at the Georgia Institute of Technology under funding from the Office of Naval Research and the Joint Deployment Agency of the Joint Chiefs of Staff introduced a new concept of interactive optimization to attack complex logistics planning problems. Interactive optimization blends the decision-making talents and instincts of experienced logistics planners with the number-crunching skills of computers. One key component is optimization-based tools, a set of prescriptive models such as linear programming which are able to evaluate large numbers of solution alternatives and make suggestions and decisions in support of the planner. Another key component of interactive optimization is color graphics, which allow man and machine to communicate efficiently.

In the early 1980's Georgia Tech researchers applied the interactive optimization concept to military deployment planning, resulting in the System for Closure Optimization Planning and Evaluation (SCOPE) prototype. SCOPE demonstrated the usefulness of combining the concepts of man-in-the-loop, graphics, and mathematical optimization tools for logistics planning applications. Many of the military's current generation of deployment planning systems which supported Operation Desert Shield are direct descendants of ideas from SCOPE, including the Military Airlift Command's ADANS system and the Military Traffic Management Command's STRADS system.

Beyond the successful application of the interactive optimization concept, two important issues remain: *flexibility* and *integration*. The flexibility issue arose as a major concern of logistics planning systems that were historically built from the ground up, or custom-built. CAPS participated in the development of several custom-built logistics software systems, including the ADANS and STRADS systems, the Postal Network Model national strategic planning system for the United States Postal Service, and the PROPHET vehicle dispatching system for Coca-Cola. While each of these systems has been fielded, experience has demonstrated that building and maintaining custom software is an inherently complex, expensive, and time-consuming effort. In addition, custom-built systems do not lend themselves to change, a troubling property for dynamic environments such as military logistics planning.

To overcome the *flexibility* issue, CAPS in 1985 started development of the *CAPS Logistics Toolkit™*, a graphics-based toolbox for addressing a variety of transportation, distribution, and logistics issues. The basic idea of the toolbox concept is that flexible software should be conformed to a logistics planning problem, rather than the problem squeezed into a rigid software system. To accomplish this, the *Toolkit* is a set of data building blocks called objects, and a set of "pre-fabricated" software tools that manage the objects. Most important, the tools and objects can be assembled in an easy way to form tailored analysis procedures.

Logistics planners can work with the *Toolkit* in an interactive graphical manner to explore or prototype ideas. More powerfully, planners can tailor the *Toolkit* to construct a custom-look application indistinguishable in look and function from a software system developed from the ground up. Because of the building block approach, a tailored *Toolkit* can be developed (or modified) in a fraction of the time and cost by logistics planners without extensive programming or computer skills.

First available for commercial use in early 1990, the microcomputer-based *Toolkit* has had a dramatic impact on the way computers are used to support logistics planning. Coca-Cola is deploying the *Toolkit* in hundreds of sites across the United States as its distribution software platform. Arthur Andersen Consulting has adopted the *Toolkit* as a software platform for logistics consulting. The United States Postal Service is using the software throughout their logistics management hierarchy, from national strategic design down through local carrier routing. Numerous other national and international organizations are applying the toolbox technology with great success.

The remaining issue, *integration*, is the focus of the CAPS SBIR effort "Interactive Distributed Logistics Management." Most of the fielded logistics planning systems are geared to planners performing the same kinds of tasks using the same levels of information. These systems do not decouple the planning decisions into the issues reflecting the actual decision hierarchy of the organization. What is needed is a methodology which integrates logistics planners at different locations and levels of the logistics management hierarchy. Such a methodology would allow different planning levels to simultaneously and cooperatively build and administer a logistics plan.

Computer-aided tools for distributed logistics planning resolve the key issue of integrating the higher level planning functions with the lower level execution-oriented functions. In the military airlift deployment problem, the lowest levels must know specifically when planes are scheduled, exactly what their load configurations are, and other detailed information such as balancing and fuel considerations. At higher planning levels, this expertise may not be present, nor is it needed in the same depth of detail as a lower level. Different types of expertise are needed at this level to make tactical decisions such as resource allocations. Distributed logistics planning methodology allows the effect of design decisions of one level to be studied on lower levels, and provides a mechanism for lower levels to filter information and needs up through the decision hierarchy. In addition, a distributed architecture is asynchronous and computationally more powerful, making practical the dynamic coordination between planning levels.

Distributed planning systems are also better suited for important security and survivability issues. Distributed planning has built-in redundancy so that the planning and execution process can continue even if sections of the planning hierarchy are not present. A distributed system also decentralizes need-to-know concerns, so sensitive issues are naturally compartmentalized according to the responsibilities of a particular planning level. A distributed planning system is also ideal for addressing what-if and resource acquisition issues, as different military strategies and resource configurations can be examined in a manner very similar to how a proposed plan actually would be executed.

2. PHASE I RESEARCH OBJECTIVES

The focus of the Phase I research was the development of a working prototype which concretely demonstrates methodology for distributed logistics planning. The key research issues were the following:

- (1) Determine appropriate solution tools and their relationships for different levels of the decision hierarchy.
- (2) Develop tools which dynamically coordinate the decision hierarchy levels as a plan unfolds, particularly between lower execution-oriented levels and higher planning-oriented levels.
- (3) Determine the appropriate selection and configuration of hardware and software tools which support distributed real-time communications and database management.

3. PHASE I ACCOMPLISHMENTS

In Phase I, we explored techniques for logistics planning in a distributed environment. The discovery process led us to identify issues common to distributed logistics systems. We believe we have devised an effective new treatment for problems which are difficult due to their size and multi-level nature. In Phase I, we identified and characterized a realistic deployment problem and applied our expertise to manage the problem in a coordinated manner never before documented. Strategic, tactical and operational issues were addressed in the coordinated effort. The most significant contributions of our model are its ability to effectively manage distributed logistics system planning by incorporating different objectives, constraints, and associated problem-solving techniques at different levels of the supervisory hierarchy and its associated use of optimization and heuristic problem-solving techniques.

3.1 Prototype Problem Scenario

A representative material (troops, equipment, food, etc.) deployment scenario was chosen as the prototype problem set. The provision and control of resources for the transport of cargo is a complex logistics problem. The coordinated planning of aircraft allocation and scheduling of cargo deployment boils down to answering the question: How many planes are needed, at what time, and how should they be used to best transport cargo?

This problem is not only topical; it is a realistic challenge to current logistics analysis methodology. Holistic solution of the airlift allocation and scheduling problem is extremely difficult due to its many dimensions. A great number of decision variables may vary independently, even in seemingly small problems. A limited resource of cargo jets, used to deploy units of cargo from origin locations to required destinations, must be spread over all theatres of operations. Deployment of cargo within the theatres is complicated by the requirement that cargo be moved only during specified time windows (due to strategic requirements, spoilage, etc.). Previous attempts to address these problems have principally focused on either the strategic (allocation) or the tactical and operational (scheduling) levels of the problem, but not both. In so doing, they failed to allow high- and low-level planners to work simultaneously in the same problem domain. These methods failed to adequately address the interdependency of low-level schedule optimization and high-level strategic allocation. The opportunity for insightful results makes the aircraft allocation and cargo deployment scheduling problem a perfect domain for our research on distributed logistics system management.

3.2 Distributed Logistics Management Approach

Our approach to the allocation planning and cargo deployment scheduling problem consists of mathematical problem definition, problem decomposition, examination of alternative constraint families and objective functions, and the examination of the use of optimization and heuristic methods for solution of the decomposed problem. This basic research was geared specifically to the description of a distributed model of logistics management and to the eventual development of a working prototype to demonstrate research results.

The underlying theoretical foundation of our model is a linear programming representation. We studied the airplane allocation and scheduling problem defined by the following.

Given:

- a set of "origins" ($i = 1, \dots, n$) and, for each origin i , a set of n_i "destinations" ($j = 1, \dots, n_i$) at known and fixed locations. A origin-destination pair is a "channel". Material arrives at each origin and must be shipped out to the destinations. Material becomes available in "loads", which are bundles of material that become available at specific times and must be delivered to a specific destination by a specific due date. The size of a load is measured in some single dimension such as weight. Loads can be split and transported

on more than one airplane.

- an integer-valued function $K(t)$ that tells how many airplanes are available at time t . All airplanes fly at the same speed and are of identical capacity c .

Question: How many planes are needed and how should they be allocated among channels over time to convey material?

We built a hierarchical control system to interactively build a near optimal solution to this problem. The levels, from highest to lowest, of this system are as follows.

1. system manager makes the gross allocations of airplanes to channels;
2. origin managers schedule airplanes on the channels emanating from their origins;
3. channel managers decide which loads to put on which scheduled flights.

Not coincidentally, our distributed model of planning and scheduling mimics the underlying management structure. Aggregation of information is a natural and desirable precursor to the decision making process when decisions must be based upon large amounts of information. In the typical organization, detail information is contained and processed at the lowest levels of the organizational hierarchy. At successively higher levels, information is combined and summarized, and decisions made at the higher levels are based only indirectly on the detail information. In accordance with the ideal of data encapsulation, it is fundamental to our approach that highly detailed information be managed and processed only where necessary and that all shared information be of the highest degree of aggregation reasonably achievable. The highly aggregated data is in essence more "portable" and easily managed.

Once during each planning period (a month, for example), the system manager will compute approximate allocations of airplanes to channels by solving a large linear program. This might be done on the Korbex system currently owned by the US Air Force or on some other machine dedicated to solving LPs. The LP solution will represent an "ideal" solution toward which the lower level managers will aim. Each origin manager will schedule airplanes to be as consistent as possible with the ideal solution passed down to him. Finally, each channel manager will solve a transportation problem (a special kind of linear program) to assign loads to scheduled flights.

Thus the system manager sketches out the shape of a solution, which becomes the goals for the lower level managers. The lower level managers use optimization based on local data to build detailed solutions close to the goal suggested by system manager.

The system manager solves the following problem, which is abstracted from the original problem by omitting identities of the loads and aggregating data about material availabilities and requirements.

Given:

- a set of "origins" ($i = 1, \dots, n$) and, for each origin i , a set of n_i "destinations" ($j = 1, \dots, n_i$) at known and fixed locations. A origin-destination pair is a "channel". Material of a single type is to be shipped from origins to destinations. Material is measured by a single dimension, such as weight.
- for each channel (i,j) the origin i has an availability function $A_{ij}(t)$ of cumulative material available for shipping to j by time t and a requirements function $R_{ij}(t)$ of cumulative material required to be shipped to j by time t . Since these functions measure cumulatives, they are nondecreasing.
- an integer-valued function $K(t)$ that tells how many airplanes are available at time t . All airplanes fly at the same speed and are of identical capacity c .

Question: How many planes are needed and how should they be allocated among channels over time to convey material?

Shipping requirements might more naturally be associated with each destination as requirements for material received. However, for simplicity we convert this into shipping requirements at the origin by translating each requirements curve back in time by an amount equal to the one-way travel time between origin and destination. Notice that this assumes that travel times are reasonably predictable; in particular, we assume that there are no "excessive" delays.

The cumulative shipment of material by time can be described by a non-decreasing "shipment function" $F_{ij}(t)$. $F_{ij}(t)$ will be a step function because material leaves in pulses (airplane loads). Note that for any time t , the average amount of material shipped per unit time is $F_{ij}(t)/t$, which is the slope of a line drawn from the origin to the point $(t, F_{ij}(t))$.

We model the constructing a shipment schedule for a destination by the finding of a nondecreasing function $F_{ij}(t)$ with the following properties: It must respect material availabilities, so that $A_{ij}(t) \geq F_{ij}(t)$; and as much as possible it should see that material is shipped in time, so that $F_{ij}(t) \geq R_{ij}(t)$. Note that the first is a hard constraint and the second is a soft constraint.

We restrict our planning to a finite horizon and discretize time into a finite number of planning periods, $t = 1, \dots, T$. We will build a linear programming based model that will allow us to compute a "rolling schedule": We will compute an optimal allocation of airplanes based on forecast availability and requirements through the planning horizon; then we will implement the allocation for the first planning period, extend the planning horizon, forecast new availabilities and requirements, and solve again.

Note that by restricting our attention to a finite set of discrete planning periods, we are throwing away some of the information in $A_{ij}(t)$, $F_{ij}(t)$, and $R_{ij}(t)$ and are implicitly treating these functions as if they were piecewise linear between planning periods. Accordingly we change notation slightly as follows. The cumulative availability of material at the beginning of planning period t is $A_{ij,t}$; the cumulative shipping requirements to be met by the end of time period t is $R_{ij,t}$; and the cumulative amount shipped by the end of planning period t is $F_{ij,t}$.

We assume that the planning horizon and planning periods have been chosen so that the following holds.

Fundamental Assumption: The length of a planning period is much greater than the round-trip travel time to any destination (say, an order of magnitude greater).

Thus, for example, it might be that the planning period is one month while round trip flight times require no more than 3-4 days. For origin i and destination j , let z_{ij} be the ratio of round-trip travel time between i and j and the length of the planning period. The results of our model will be more accurate as the largest z_{ij} decreases (that is, as length of the planning period increases). The length of the planning period should probably be at least 10 times greater than the largest round-trip travel time between any origin and destination.

Because of discretized time and the Fundamental Assumption, we are imputing essentially instantaneous delivery of material from each origin to any destination it serves; that is, any flight that departs the origin during planning period t is assumed to arrive at its destination before the beginning of planning period $t+1$. Again, this will be more accurate as the largest z_{ij} decreases (that is, as length of the planning period increases with respect to the largest round-trip travel time).

Now we develop a linear program that represents the problem of allocating airplanes to channels.

The constraints that material cannot be shipped from origin i to destination j until it is available at i can be represented as follows.

$$F_{ijt} \leq A_{ijt} \quad (\text{cumulative availability})$$

For convenience we convert this to constraints on availability within each planning period. First add a slack variable to each constraint so that they become

$$F_{ijt} + I_{ijt}^1 = A_{ijt}$$

The slack variable I_{ijt}^1 can be interpreted as the inventory intended for destination j that is carried at origin i during planning period t . Now subtract the constraint for $i, j, t-1$ from the constraint for i, j, t and write the amount shipped during (not "by") planning period t as f_{ijt} and amount that becomes available during planning period t as a_{ijt} .

$$f_{ijt} - I_{ijt-1}^1 + I_{ijt}^1 = a_{ijt} \quad (\text{incremental availability}).$$

Similarly, the constraints that shipping requirements must be met can be represented as

$$F_{ijt} \geq R_{ijt} \quad (\text{all } i, j, t) \quad (\text{cumulative shipping requirements}).$$

Again we convert this to constraints within each planning period. First add a surplus variable to each constraint so that they become

$$F_{ijt} - I_{ijt}^2 = R_{ijt} \quad (\text{all } i, j, t).$$

The surplus variable I_{ijt}^2 can be interpreted as the inventory carried at destination j during planning period t that was delivered from origin i in advance of its need. Now subtract the constraint for $i, j, t-1$ from the constraint for i, j, t and write the amount shipped during (not "by") planning period t as f_{ijt} and amount that is required during planning period t as r_{ijt} .

$$f_{ijt} + I_{ijt-1}^2 - I_{ijt}^2 = r_{ijt} \quad (\text{incremental shipping requirements}).$$

In addition, for each planning period there is a constraint that limits the number of airplanes available. We derive that constraint as follows. During each time period f_{ijt} gives the amount of material shipped by planning period t . Since each plane is of capacity c , the number of flights required is f_{ijt}/c . This can be accomplished with approximately $P_{ijt} = f_{ijt} z_{ij}/c$ airplanes. (Note that we are here allowing fractional airplanes; this will be resolved later.) Therefore the constraints on airplane availability become

$$P_{ijt} - f_{ijt} z_{ij}/c = 0 \quad (\text{planes on channel } (i,j) \text{ during period } t),$$

$$\sum_i P_{ijt} \leq K_j \quad (\text{airplane availability during period } t).$$

In summary, the following sets of constraints define the restrictions on how material can be shipped.

$$f_{ijt} - I_{ijt-1}^1 + I_{ijt}^1 = a_{ijt} \quad (\text{all } i, j, t) \quad (\text{incremental availability}),$$

$$f_{ijt} + I_{ijt-1}^2 - I_{ijt}^2 = r_{ijt} \quad (\text{all } i, j, t) \quad (\text{incremental requirements}),$$

$$\sum_j I_{ijt}^1 \leq K_i \quad (\text{all } i, t) \quad (\text{inventory capacity at origin } i \text{ during } t),$$

$$\sum_i I_{ijt}^2 \leq K_j \quad (\text{all } j, t) \quad (\text{inventory capacity at destination } j \text{ during } t),$$

$$P_{ijt} - f_{ijt} z_{ijt}/c = 0 \quad (\text{all } i,j,t) \quad (\text{planes on channel } (i,j) \text{ during period } t),$$

$$\sum_t P_{ijt} \leq K_i \quad (\text{all } i) \quad (\text{airplane availability during period } t).$$

One of the advantages of this linear programming model is that there are many possible useful objective functions, including the following.

- simple feasibility: Are there enough airplanes available during each planning period to ensure aggregate feasibility?

- minimize the number of airplanes in the fleet required to ensure aggregate feasibility: All $K_i = K$ and the objective is to minimize K .

- minimize total inventory costs: Material incurs inventory costs wherever it is held before use, either at the origin from which it is shipped or at the destination if shipped in advance of use. Let c_{it} , c_{jt} be the costs of holding one unit of inventory at origin i or destination j , respectively, during planning period t ; then to minimize total inventory costs is to minimize $\sum_t [c_{it} I_{it} + c_{jt} I_{jt}]$.

The basic linear programming model can be extended in many ways. Some useful extensions are: the allowance of planning periods of arbitrary and possibly different lengths; the deferral of requirements at some cost; inclusion in the linear model of plane transport to new channels in successive planning periods; the allowance that destination requirements can be met from multiple origins; and the allowance of airplanes of differing capacities.

There is no requirement in the current formulation that all planning periods be the same length. One can use smaller planning periods to examine critical periods more closely and use longer planning periods toward the end of the planning horizon, when data are less reliable. The LP formulation is independent of the durations of the planning periods and need not be changed.

Deferred requirements can be implemented by letting B_{ijt} be the amount short at destination j from origin i during planning period t and augmenting the constraints on incremental requirements to become

$$f_{ijt} + I_{ijt} - I_{ijt} - B_{ijt} + B_{i,j,t+1} = r_{ijt} \quad (\text{all } i, j, t)$$

Also include total "backorder" costs $\sum_{ijt} b_{ijt} B_{ijt}$ in the objective function.

In general, each channel will have different numbers of airplanes assigned to it in different planning periods and from one planning period to another airplanes will have to be reallocated among channels. An approximate way to do this is to solve a transportation problem among the channels. This might be adequate if the transportation costs to reallocate planes is small compared to the cost of operating them at a channel. If, however, transportation costs are significant, then the transportation problem can be embedded in the allocation program to trade-off allocation costs and transportation costs optimally. To do this, let $P_{ijk,t}$ be the number of planes transported to channel (i,j) from channel (k,l) at the beginning of planning period t . Then conservation of planes requires the following.

$$P_{ijt} = \sum_k P_{ijk,t} \quad (\text{all } i, j, t) \quad (\text{conservation of airplanes})$$

Let $c_{ijk,t}$ be the cost of transporting a plane to channel (i,j) from channel (k,l) at the beginning of planning period t . Then augment the objective function with the total transportation costs $\sum_{ijk,t} c_{ijk,t} P_{ijk,t}$.

In the case where destination requirements can be met from multiple origins, we can rewrite a_{jt} as a_{jt} and r_{ijt} as r_{jt} . Now it must be that total shipments to destination j from all origins meets requirements, which can be expressed as follows.

$$\sum_j f_{jt} - I_{i,t-1} + I_{it} = a_{it} \quad (\text{all } i, j, t) \quad (\text{incremental availability})$$

$$\sum_i f_{jt} + \sum_i I_{j,t-1} - \sum_i I_{jt} = r_{jt} \quad (\text{all } j, t)$$

Extension of the prototype to allow for differing airplanes is achieved as follows. We assume that all planes travel at roughly the same speed so that travel times remain generally unchanged. Therefore the only significant difference in our model is that each plane type has a different capacity. For example, let there be two types of planes, with capacities $c1$ and $c2$ respectively. Then the total material sent from origin i to destination j during planning period t can be expressed as the sum of material sent on planes of type 1 and material sent on planes of type 2:

$$f_{jt} = f_{jt}^1 + f_{jt}^2.$$

Then the number of planes required on each of these channels during this planning period are as follows.

$$P_{jt}^1 = f_{jt}^1 / c1 = 0 \quad (\text{type 1 planes on channel (i,j) during period t),}$$

$$P_{jt}^2 = f_{jt}^2 / c2 = 0 \quad (\text{type 2 planes on channel (i,j) during period t).}$$

Finally, the allocation of planes must respect availability of the various types during the planning period.

$$\sum_j P_{jt}^1 \leq K_t^1 \quad (\text{availability of type 1 planes during period t),}$$

$$\sum_j P_{jt}^2 \leq K_t^2 \quad (\text{availability of type 2 planes during period t).}$$

Note that a possible problem with enlarging the model this way is that in the LP solution the P_{jk} will be smaller and therefore more likely to be difficult to interpret because of fractional values.

The LP model looks at the logistics system only at discrete points in time and assumes that the system is "well-behaved" (approximately linear) during the periods between those points. Thus, in effect, the availabilities and requirements are being modelled as piecewise linear over time.

The LP does not require uniform planning periods and therefore the manager can select planning periods as desired. This allows adaptive modelling in which, for example, the planning periods can be chosen to be of shorter duration for increased accuracy when the data of the problem (availabilities and requirements) are changing quickly. On the other hand, when the data are changing only slowly, then one can use long planning periods, which will help keep the model small and manageable.

There appears to be, in some sense, a "best" selection of planning points. If the planning periods are all too short, then the LP grows to possibly impractical size; but there is also the problem that it loses accuracy as the Fundamental Assumption is increasingly weakened. On the other hand, if planning periods are too long, then the model becomes inaccurate if the availabilities and requirements are other than linear during planning periods. Since the second type of inaccuracy seems less severe, this suggests that managers err on the side of choosing long planning periods. In any event, one should use time periods that are as long as possible while still modelling the availabilities and requirements.

It might be useful to present the managers with graphs of availabilities and requirements and allow them to select planning periods by pointing and clicking. Then they could interactively select the appropriate level of detail at which to examine the logistics system. Meanwhile, it remains to be studied how to compute the best choice of planning points.

The LP model above will allocate fractional amounts of airplanes and this must be reconciled when the LP solution is implemented as schedules for airplanes. Assume the LP to have allocated $\{P_{jt}\}$ airplanes among the channels $\{i,j\}$. Then $\text{FLOOR}(P_{jt})$ planes can be devoted exclusively to channel $\{i,j\}$ during t , leaving $\text{FRAC}(P_{jt}) = P_{jt} - \text{FLOOR}(P_{jt})$ "notional planes" to be accounted for. For convenience we write $p_{jt} = \text{FRAC}(P_{jt})$. The lower level

managers must decide how and whether to provide the $\{p_{ijt}\}$ notional planes. The simplest but least accurate way of dealing with this is to simply round any allocation down, so that, while the LP allocated P_{ijt} to channel (i,j) , only $\text{FLOOR}(P_{ijt})$ airplanes are actually sent. This, however, will tend to underequip channels and so tend to fail to meet shipping requirements. It would also leave unused, excess planes whose number would be on the order of the number of channels.

A similarly simple but inaccurate strategy is to overequip each channel by simply rounding up airplane assignments so that channel (i,j) is allotted $\text{CEILING}(P_{ijt})$ airplanes. This has the advantage of providing each channel with some excess plane capacity to protect against unforeseen developments. However, it also requires additional airplanes beyond those presumably available. In fact, the number of additional airplanes can be on the order of the number of channels in the distribution system.

A more effective way of providing the notional airplanes is for the origin to share airplanes among its channels. For example origin i could share $\text{CEILING}(\sum_j p_{ijt})$ notional airplanes among its channels so that there would be one of the notional airplanes flying channel (i,j) about fraction p_{ijt} of the time.

It is an interesting problem to compute a strategy for sharing airplanes among channels of a common origin. Suppose, for example, that the $p_{ijt} = 0.6, 0.3, 0.3, 0.2$ for $j = 1, \dots, 4$. Then there are $\text{CEILING}(0.6 + 0.3 + 0.3 + 0.2) = 2$ notional airplanes required at origin i , which is more than the 1.4 allocated by the LP. Now we have to determine exactly how to share the 2 notional airplanes. One way is to assign one notional plane to fly about 60% of the time on channel $(i,1)$ and about 30% of the time on channel $(i,2)$; and assign another notional plane to fly about 30% of the time on channel $(i,3)$ and about 20% of the time on channel $(i,4)$ (with some additional time off). Another implementation would be to assign one plane to fly channel $(i,1)$ 60% of the time and channel $(i,2)$ 10% of the time; and assign the second plane to fly channel $(i,2)$ 20% of the time, channel $(i,3)$ 30% of the time, and channel $(i,4)$ 20% of the time. In any case we will have exceeded the LP solution by $2 - 1.4 = 0.6$ planes. However--and this is the improvement over the straightforward technique--the number of additional airplanes required beyond the LP solution will be in proportion only to the number of origins and not to the number of channels.

Even this overcommitment of airplanes can be avoided by instituting sharing among origins as well as sharing among channels. For example, each origin i could coordinate sharing among channels of only $\text{FLOOR}(\sum_j p_{ijt})$ of its notional airplanes; the remaining $\text{FRAC}(\sum_j p_{ijt})$ notional airplanes assigned to origin i could be controlled by the higher level manager, who would coordinate sharing of the remaining notional airplanes, which, summing over all origins, are $\sum_i \text{FRAC}(\sum_j p_{ijt})$ in number, among all origins (and so no more than the number of origins). The advantage of this multi-level sharing is that it would not require any airplanes beyond those given in the availability constraints (5) and (6). It is worth remarking that even when we know that a notional airplane should fly channel (i,j) about fraction p_{ijt} of the time, there remains the question of the details of his schedule. This could be left to the discretion of the manager at the origin, who might, for example, choose simply to dedicate that airplane to channel (i,j) for the first $(p_{ijt} \times 100)$ percent of the planning period, and afterwards reallocate the plane to another channel. Alternatively, the manager at the origin might prefer to schedule via a heuristic that will tend to share the plane in the desired proportions. One such heuristic is the following:

To schedule a notional airplane who must service channels (i,j) , $j = 1, \dots, n_i$, according to the fractional allotments $\{p_{ijt}\}$:

- Each time an airplane returns to the origin, choose the next channel to fly to be one for which the difference between ideal allocation and actual allocation so far is maximum.

This heuristic has the useful property that it shares the airplanes according to their ideal proportions as the largest z_{ij} decreases (the ratio of round-trip travel time between i and j and the length of the planning period).

A similar strategy could be used by the higher level manager to coordinate sharing of notional airplanes among the origins.

At the operational planning level, let there be L loads to be delivered during the first time period. Each load l is distinguished by its arrival time a_l at the origin and its due date d_l at the destination.

Let $C(t)$ be the cargo-carrying capacity available at time t . Once the schedule of airplanes is known, then $C(t)$ is known.

The channel manager operates according to smaller planning periods than do the higher level managers. We refer to these as planning subperiods; let them be T' in number. For example, while the system manager will plan by the month, for example, the channel manager might plan by the day or even by the hour.

We approximate $C(t)$ by the series $\{C_t\}$ ($t = 1, \dots, T'$). Now consider the following graph. There are L vertices on the left, each representing a load to be delivered; and there are T' vertices on the right and each represents a planning subperiod. There is an arc between load l and each planning subperiod a_l, a_{l+1}, \dots, d_l . There are also arcs from each planning subperiod to a "sink" and the capacity of each is the total cargo-carrying capacity of all airplanes departing on that channel during that planning subperiod. Finally, each vertex corresponding to a load produces "flow" in the network equal to the size of the corresponding load. Now if we maximize the flow from the left of the graph to the sink, we will determine an assignment of loads to airplanes that minimizes the shortfall (that is, the amount of material not delivered by its due date). In fact we can maximize the flow to the sink very quickly by appealing to the special structure of the network: It is "convex" because if there are arcs from l to t_i and from l to t_j ($t_i \leq t_j$), then there must be arcs from l to t_k for all $t_i \leq t_k \leq t_j$. This flow problem can be solved optimally by a single-pass algorithm that operates as follows: Load the next-departing airplane with those available loads that have the earliest due date.

3.3 Model of Distributed Logistics Management

Based on the preceding theoretical exposition, our model is a two-level system incorporating a single senior planner and multiple junior planners, all linked in a dynamic environment. The senior planner is responsible for the strategic allocation of airlift resources to meet the requirements of the junior planners. The junior planners are responsible for the tactical scheduling of their airlift allocations and the operational loading of the scheduled aircraft. The nominal objective of the model is the minimization of required airlift resources subject to the on-time fulfillment of all movement requirements. It is important to note, however, that the focus of the research was not to find the best way to solve the prototype problem; instead, the intent was to demonstrate methodology for *distributed* logistics planning given a *plausible* way to solve the problem.

The planners at each level work simultaneously and independently of each other, but in concert with their workstation microcomputers, to realize their objectives. At the strategic planning level, the plan objective is to satisfy the aggregate demands for aircraft of each junior planner, while observing limitations on total aircraft resource available. At the tactical/operational planning level, the plan objective is the satisfaction of movement requirements, subject to the aircraft capacity supplied by the high-level planner.

The junior planners work with detailed data to (1) transmit resource requirements up to the senior planner and (2) realize the aggregate plan obtained from the senior planner. The senior planner works with aggregate information transmitted from the junior planners to address high-level design issues in developing the aggregate plan. Information transmission is conducted instantaneously via a local area network and a dedicated database machine which handles information transactions invoked by planning workstations.

The system encompasses, top to bottom, (1) system-wide allocation of airlift resources among service areas (generation of allocation functions for each channel), (2) scheduling of aircraft departures within services areas and (3) loading of cargo onto the scheduled aircraft.

4. PROTOTYPE IMPLEMENTATION AND FUNCTIONAL DESCRIPTION

4.1 Overview

The prototype is a small-scale realization of the two-level system described above. It incorporates a single senior, strategic planner and two junior, tactical/operational planners, all linked via a local area network. The implemented objective of the prototype is the minimization of required airlift resources subject to the on-time fulfillment of all movement requirements. Several data paths are implemented for the exchange of aggregate data between the planning levels.

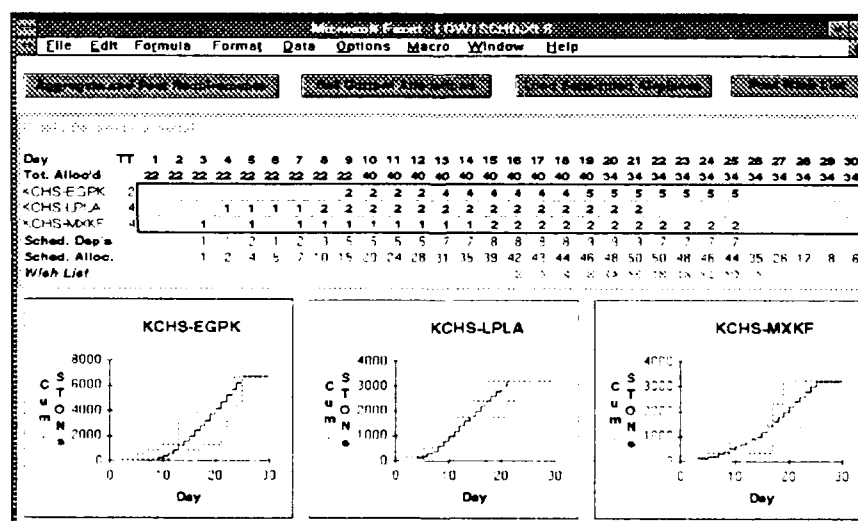
It was the purpose of our Phase I research to prototype methodology for distributed logistics system management. The purpose was not to develop production-level tools at this juncture, and neither the data exchange devices nor the planning modules are capable of withstanding the increased load of general use. However, the prototype is functional and effective. Our focus during Phase I was to perform basic research into distributed logistics planning. It will be our focus during Phase II to conduct applied research and development in the distributed logistics management arena, including the development of powerful, robust tools for deployment in our flagship logistical analysis product, *CAPS Logistics Toolkit*.

The following sections detail the functions of the low-level tactical/operational planning modules, the high-level strategic planning module and the information management system.

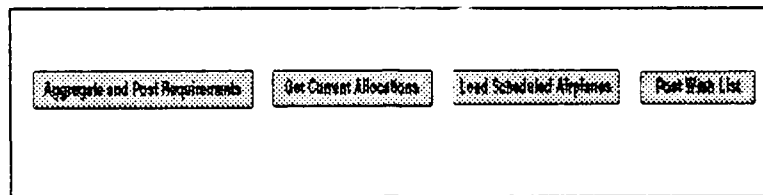
4.2 Low-Level Tactical/Operational Planning Modules

The low-level modules have both their front ends and computational engines in Microsoft Excel. They incorporate multiple worksheet, graph and macro files which work in tandem to perform both heuristic and optimal algorithmic tasks. At the low level, the junior planner, acting through the Microsoft Excel interface, is responsible for requirements aggregation and posting, airplane scheduling, and wish list posting for a fixed (30 day) time horizon.

All actions of the junior planner are performed using the interface shown below. The low-level interface consists of three regions (top to bottom): the button region; the schedule table region; and the graphical feedback region. The regions are differentiated not only by location, but also by the objects from which they are composed.



The button region contains the four labeled buttons shown below.



Clicking the mouse on any of the buttons invokes its linked macro program which automatically perform repetitive and sometimes complicated tasks. Operations available from buttons include: (1) aggregation of MRs and posting of the aggregate requirements to the information server; (2) extraction of current airlift allocations from the server; (3) optimized airplane loading for the current airlift schedule; and (4) posting of the current wish list to the server.

Clicking the mouse on the button labeled "Aggregate and Post Requirements" invokes a command macro which "rolls up" the individual MRs into the a_{ijt} and r_{ijt} availability and requirement functions explained previously. MR aggregation encompasses, bottom to top: (1) the categorization of MRs by channel; (2) the calculation of availability and requirements functions for each MR; and (3) the aggregation of the MR functions into availability and requirements functions a_{ijt} and r_{ijt} for the MRs of each unique channel. Requirements posting by the junior planner consists of "publishing" the computed a_{ijt} and r_{ijt} functions by posting them to the information server.

Clicking the mouse on the button labeled "Get Current Allocations" invokes a command macro which extracts the current allocations for the low-level planner's channels from the information server. The total allocation of airlift to the origin by day is placed in the schedule table on the row labeled "Tot. Alloc'd."

Clicking the mouse on the button labeled "Load Scheduled Airplanes" invokes an optimal procedure which determines the best loading of MRs onto the airlift schedule currently imposed on the schedule table. An information window is displayed at algorithm completion which details the percentage of MRs moved on time both in total and by channel.

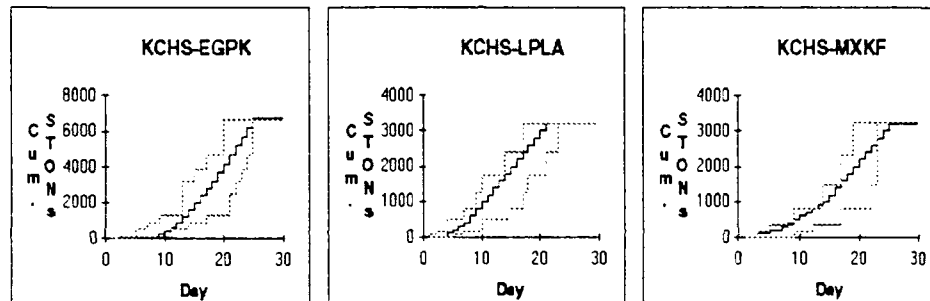
Clicking the mouse on the button labeled "Post Wish List" causes the current wish list vector w_{it} (labeled "Wish List" in the schedule table region) to be posted to the information server. As previously noted, only one wish list vector is defined for each origin, not each channel.

The schedule table region of the interface consists of the textual object shown below.

Day	TT	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Tot. Alloc'd		22	22	22	22	22	22	22	22	22	22	40	40	40	40	40	40	40	40	34	34	34	34	34	34	34	34	34	34	34	34
KCHS-EGPK	2										2	2	2	2	4	4	4	4	4	5	5	5	5	5	5	5	5	5	5	5	5
KCHS-LPLA	4				1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
KCHS-MXXF	4			1	1	1	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2	2
Sched. Dep's			1	2	1	2	3	5	5	5	5	7	7	8	8	3	3	9	9	3	7	7	7	7	7	7	7	7	7	7	7
Sched. Alloc.			2	4	5	7	10	15	20	24	28	31	35	39	42	43	44	46	48	50	50	48	46	44	35	25	17	8	5	5	
Wish List																															

The central region of the interface is the schedule table. This table is the centerpiece of an effective interactive airlift scheduler. Onto this table the number of aircraft departures from the origin by day, by channel are specified for the entire planning horizon. Days are numbered horizontally along the top of the table. Channels, as well as summary information, are listed down the left side of the table. The "Tot. Alloc'd" row, as noted above, contains the current origin airlift allocation extracted from the information server. The "KCHS-EGPK," "KCHS-LPLA" and "KCHS-MXKF" channel labels indicate that the information contained in the corresponding rows is relevant only to the named channel. The "Sched. Dep's" row contains the calculated requirement of aircraft departures for the specified day, based on the current departure schedule. The "Sched. Alloc." row contains the calculated requirement of airlift for the specified day, once again based on the current departure schedule. For the calculation of required airlift on specified days, round trip travel and immediate return trip is assumed on all channels.

Textual summary information is not the only dynamically-calculated feedback information supplied by the low-level system. Under the schedule table is the graphical feedback region, shown below. In this region the system provides a graphical representation of the cumulative availability, requirement and allocation functions A_{ijt} , R_{ijt} and F_{ijt} functions (cumulative a_{ijt} , r_{ijt} and f_{ijt} respectively) for each channel. These three graphs are continuously updated to reflect the departure schedule and requirements on each channel.



But the tactical/operation system *as a whole* is much greater than the sum of its parts. The integration of the three interface regions gives the planner instantaneous and powerful feedback based on complex resource requirements, changing allocations and interactive schedule choices as he works to best realize his airlift allocation. In summary, his task consists of the scheduling of airplane departures and the subsequent loading of cargo onto scheduled aircraft.

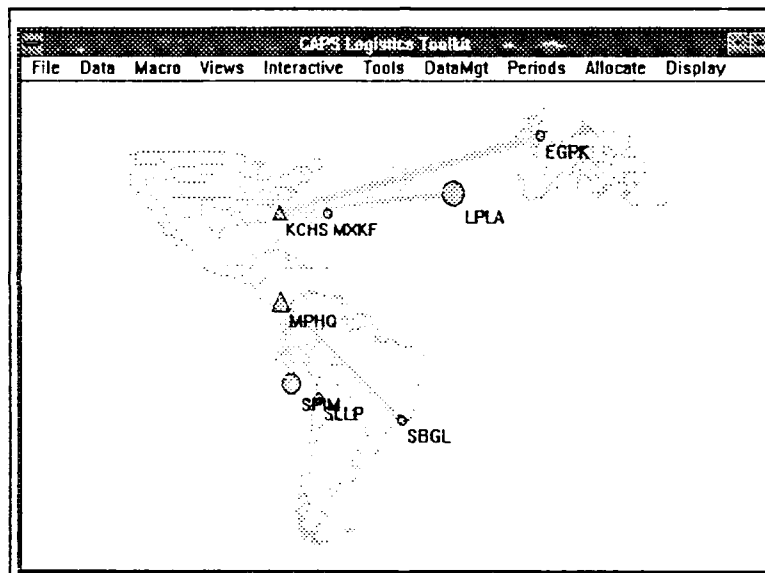
The process of departure scheduling is performed manually at the junior planner's interface by "laying in" departures in the tabular interface shown in the diagram. Beneath the input table, the interface produces a continuously updated tabular output of resources used versus resources allocated as well as a continuously updated graphical representation of airlift applied versus required for each channel. The planner is expected to lay in a schedule which (approximately) fits its F_{ijt} allocation curve between the availability A_{ijt} and requirement R_{ijt} curves for each channel, while respecting the limits of aircraft allocation to the origin of the channels. Scheduling of aircraft which are not allocated is allowed and results in one or more entries in the "Wish List" row of the tabular interface. Wish list values indicate additional planes desired beyond the current allocation. Note that an F_{ijt} curve which falls below the R_{ijt} curve represents insufficient airplane capacity, and an F_{ijt} curve which exceeds the A_{ijt} curve represents wasted capacity.

When the junior planner gets a schedule that fits the aggregate requirements, he can quickly determine its effectiveness with the actual detail data by clicking the mouse on the "Load Scheduled Airplanes" button to *optimally* load cargo units according to the current schedule. Finally, when a "best" schedule has been created, wish list information can be posted to the information server by clicking the mouse on the "Post Wish List" button.

4.3 High-Level Tactical Planning Module

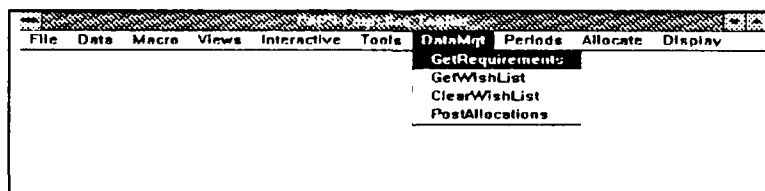
The high-level system is designed on the CAPS Logistics Toolkit system. The Toolkit has facilities for network flow model (NFM) solution, color graphics for network and geographic map display and a highly developed macro programming language. Exploiting such capabilities, the high level system allows custom-configurable display of the allocation of resources among channels of the prototype system. At the high level, the senior planner is responsible for the extraction of current requirements from the information server, allocation of aircraft to the origin sets for a fixed (30-day) time horizon and posting of allocations to the information server.

The fundamental display of the high-level system is shown below. It consists of a geographic map overlaid with the origins, destinations and channels of the distributed logistics system.



All functionality of the high-level system is available through a custom-configured Toolkit menu. Selection, by mouse click, of any of the rightmost four menu items "DataMgt," "PlanningPds," "Allocate," and "Display" opens a corresponding pull-down menu from which the functions of the tactical planning module can be accessed. The following sections detail the functionality available from each of the menu selections.

Selection of the "DataMgt" item opens the data management pull-down menu shown below.



The data management pull-down menu provides options of extracting the aggregate resource requirements from the information server, extracting the wish lists from the server, clearing the wish lists from the allocation networks, and posting new airlift allocations to the server.

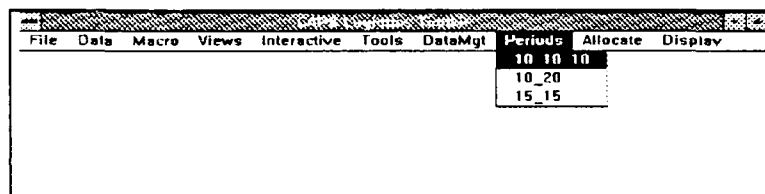
Selection of the "GetRequirements" option invokes a Toolkit macro program which extracts the current aggregate resource availability a_{ijt} and requirement r_{ijt} from the information server. The macro causes the network structures local to the strategic planning workstation to be updated with the current requirement vectors from the server.

Selection of the "GetWishList" option invokes a macro which extracts the current wish lists w_{ijt} from the information server. Similar to "GetRequirements," the macro causes the local network structures to be replaced with the current wish list information from the server.

Selection of the "ClearWishList" option invokes a macro which zeroes the wish list vectors w_{ijt} on the allocation networks. This has the effect of nullifying any wish list information so that the allocation algorithms consider only the availability a_{ijt} and requirement r_{ijt} .

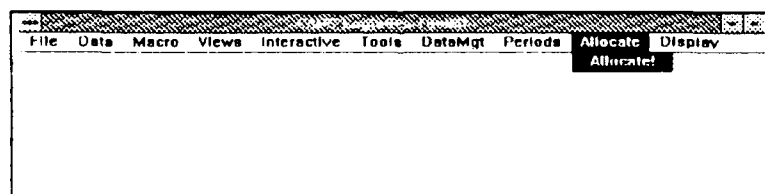
Selection of the "PostAllocations" option invokes a macro which posts the current airlift allocation to the information server, replacing the next most recent allocation. This option is invoked only after an allocation acceptable to the high-level planner is achieved.

Selection of the "PlanningPds" item opens the planning period configuration pull-down menu shown below.



The planning period pull-down menu provides three planning period options. Selection of the "10_10_10" option configures the allocation networks to produce solutions which have constant airlift allocations to origins during the first, second and third ten-day planning periods. Similarly, the "10_20" option configures the networks for solutions which have constant allocations during days one through ten and constant, but possibly different, allocations during days 11 through 30. The "15_15" option follows the same pattern, and other configurations are possible, requiring only the development of a new configuration macro and menu specification.

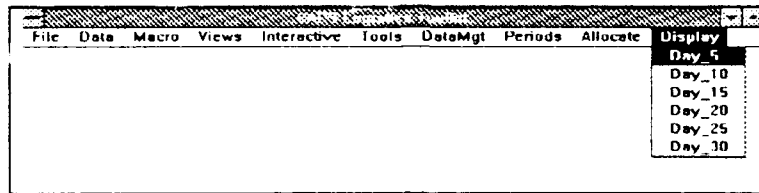
Selection of the "Allocate" item opens the allocation macro program pull-down menu shown below.



The only option available from the "Allocate" pull-down menu is "Allocate!" Selection of "Allocate!" invokes the NFM solution macro program which uses powerful Toolkit tools to reach a solution based on the pre-configured NFMs. The allocation of notional planes is conducted by solution of the NFMs using the Toolkit proprietary GeneralizedMinCostFlow algorithm. The algorithm determines the notional planes required for each channel to just satisfy the aggregate resource requirements. The allocations of notional planes to channels are summed over common

origins and rounded up to the next highest integers for each day of each planning period to produce the origin allocations.

Selection of the "Display" item opens the display options pull-down menu shown below.



From the "Display" pull-down menu, the notional plane allocations can be viewed for a specific day of the planning horizon. Pull-down menu options consist of "Day_5" through "Day_30" views in five-day increments. Selection, for instance, of "Day_5" invokes a macro which scales channel widths according to their notional plane allocation on day five of the planning horizon.

Like the tactical/operational planning module, the strategic planning module as a whole is somewhat greater than the sum of its parts. The configurable Toolkit structure and macro language make for effortless "what-if" analysis, because repeated solution of the allocation NFM's under different planning period configurations can be performed by executing a few keystrokes. Whenever an allocation plan is determined to be acceptable to the senior planner or "better than the previously posted plan" it can be transmitted to the information server by selection of the "PostAllocations" option, thus updating the plan on the server.

4.4 Information Management System

It is the management of information, both within and between workstations, that is the prototype's real contribution to innovation. Under the blanket title of information management in the prototype are Dynamic Data Exchange (DDE) and transaction-based remote data management. The pieces of the information management system work behind the scenes to lend the power of integration to the distributed logistics management system prototype.

Through DDE, applications running concurrently in the Microsoft Windows environment may share access to each other's information. The prototype system uses links at both the senior and junior planners' workstations. At the high level, links between the Toolkit and Pioneer Software's Q+E database query and update interface are used along with the Windows Clipboard to extract the availability, requirements and wish list vectors from, and to update the allocation functions to, SQL Server. At the low level, links between Excel and Q+E are used to post the availability, requirements and wish list functions to, and to extract the allocation functions from, SQL Server. In both cases, the links are "cold," meaning that local data is only updated when explicitly specified (as is the case when macro programs are run in either the Toolkit or Excel).

Transaction-based remote data management is achieved by use of a dedicated workstation tied to a Novel local area network (LAN). The dedicated workstation serves as a database engine and data repository, running Microsoft SQL Server Structured Query Language database system. Use of SQL Server not only gives shared access to data; it also increases data security, data integrity and system modularity--for as any production system grows or changes, any migration from personal computer database management to mini or mainframe management would be transparent between SQL standard servers.

Data contained on SQL Server may have highly restricted access or may be nearly public. The data may be extracted only by system entities with certain specified permissions and may be changed only by the *transactions* of system entities with the required permissions. Between transactions, the data is static. It is permanent until

explicitly changed by a party with permission to do so.

All central information for the prototype is either (1) updated by the high-level workstation but only read by the low-level workstations or (2) updated by one of the low-level workstations but only read by the high-level workstations. An entity may only post to *or* read from one of the central tables, not both. The central database consist of the *profiles* table, the *allocations* table and the *wishlist* table. The profiles and wishlist tables are written by the low-level workstations to be read by the high. The allocations table is written by the high-level workstation to be read by the low.

5. RELATIONSHIP TO FUTURE RESEARCH

Phase II will continue the research and development effort initiated in Phase I toward the development of a production distributed logistics management system. We will extend the methodology to address a broader range of distributed environment applications. Phase III efforts will focus on the realization of basic and applied research of Phases I and II. In this final phase, robust, versatile tools will be developed and deployed on our flagship product, CAPS Logistics Toolkit.

6. CONCLUSION

In Phase I, we explored and analyzed various approaches, techniques and models of distributed logistics system management. An extensive and diversely extensible model was developed for the problem of airlift allocation and scheduling subject to cargo movement requirements. A model of high-level allocation of airlift resources and low-level schedule realization was devised and extended for several constraint sets and objective functions. A working prototype has demonstrated the value of the distributed concept for the solution of the multi-level problem by incorporating the different views of strategic and tactical/operational planners. Since our distributed model closely resembles the organizational decision hierarchy for which it is designed, a realization of it will fit very naturally into existing distributed planning environments.

Also noteworthy is the behind-the-scenes power demonstrated by macro-driven information exchange between applications and machines using Dynamic Data Exchange and remote database server technology. These and other important technologies and issues for distributed management were identified and explored. Finally, research directions for Phase II research were identified.

The results of Phase I research and prototyping are very encouraging. It is our belief that Phase II and III research and design can produce robust tools for distributed logistics system management. We anticipate that deployment of these tools on the CAPS Logistics Toolkit platform will benefit both government and civilian users.